What are the fundamental structures of concurrency? We still don't know!

Samson Abramsky [1,2]

Oxford University Computing Laboratory
Oxford, U.K.

arXiv:1401.4973v1 [cs.LO] 20 Jan 2014

Abstract

Process algebra has been successful in many ways; but we don't yet see the lineaments of a fundamental theory. Some fleeting glimpses are sought from Petri Nets, physics and geometry.

Key words: Concurrency, process algebra, Petri nets, geometry, quantum information and computation.

1 Process Calculi as Generic Theories

What counts as a successful theory in Computer Science? Consider obvious exemplars such as

- Process Calculi
- Type Systems
- Model-checking

It is not the case that there is a single agreed model, notation, formalism, tool or language in any of the above areas. In fact there are a profusion of all of these, although some have been particularly influential. (Insert your favourite examples here ...)

The 'Next 700 ⋯' syndrome

Is this profusion a 'scandal' of our subject? I used to think so — and I wasn't alone (e.g. Robin Milner's quest to find the 'λ-calculus of concurrency'). Now I am not so sure.

[1] This research was supported by U.K. EPSRC
[2] Email: samson@comlab.ox.ac.uk

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

It's the *Paradigms!*

The paradigms and tool-kits, both technical and conceptual, provided by these theories have been deeply absorbed by the research communities and have increasingly influenced applications.

Examples:

- labelled transition systems and bisimulation
- naming and scope restriction and extrusion
- the automata-theoretic paradigm for model-checking
- the type systems paradigm, with compositional typing rules for terms-in-context, and key structural properties such as Subject Reduction.

By their fruits shall ye know them.

These tool-kits are *the real fruits of these theories*. They may be compared to the traditional tool-kits of physics and engineering: Differential Equations, Laplace and Fourier Transforms, Numerical Linear Algebra, etc.

They can be applied to a wide range of situations, going well beyond those originally envisaged, e.g. Security, Computational Biology, Quantum Computing, etc. So, is everything in the garden rosy?

Dreams of Final Theories

But can we do better than this? After all, in physics there *are* great theories which transcend mere tool-kits. We largely lack such theories, in Computer Science as a whole, and in concurrency and process calculus in particular. Is this unavoidable, as part of the nature of our subject, or will such theories emerge?

Some may find such questions uninteresting, or even meaningless; they can safely stop reading here.

2 Process Calculi vs. Concurrency Theory

1980 marked the start of a new era in concurrency theory, but not its beginning. A meaningful theory of concurrency, incorporating some profound insights, had been originated by Petri in the 1960's, and Net theory, as well as other approaches to concurrency, continues to be actively developed.

There is no doubt that the advent of algebraic process calculi marked a decisive advance in concurrency theory, in particular in the use of compositional algebraic methods for the description of complex systems. It is often the case, though, that when an advance is made, something valuable is also lost, or at least, temporarily forgotten.

Let us start with the problem of *canonicity* — the 'next 700 process algebras' syndrome. In a sense, the very success of the paradigmatic tool-kit, as described in the previous section, is also the source of the problem. It is too

easy to cook up yet another variant process calculus or algebra; there are too few constraints. This *plasticity of definitions* has become so familiar in our field that we may not be aware of it as an issue. The mathematician André Weil apparently compared finding the right definitions in algebraic number theory — which was like carving adamantine rock — to making definitions in the theory of uniform spaces (which he founded), which was like sculpting with snow. In concurrency theory, we are very much at the snow-sculpture end of the spectrum. We lack the kind of external reality, whether it comes from fundamental mathematical objects like the integers, or manifolds, or differential equations, or from physical reality as determined by experiment, which is hard and obdurate, and resistant to our definitions. Is this a necessary feature of our existence, or have we just not yet found the real bedrock?

An important quality of Petri's conception of concurrency is that it *does* seek to determine fundamental concepts: causality, concurrency, process, etc. in a syntax-independent fashion. Another important point, which may originally have seemed merely eccentric, but now looks rather ahead of its time, is the extent to which Petri's thinking was explicitly influenced by physics (see e.g. [6]. As one example, note that K-density comes from one of Carnap's axiomatizations of relativity). To a large extent, and by design, Net Theory can be seen as a kind of *discrete physics*: **lines** are time-like causal flows, **cuts** are space-like regions, **process unfoldings** of a **marked net** are like the solution trajectories of a differential equation.

This acquires new significance today, when the consequences of the idea that 'Information is physical' are being explored in the rapidly developing field of quantum informatics. Moreover, the need to recognize the spatial structure of distributed systems has become apparent, and is made explicit in formalisms such as the Ambient calculus, and Milner's bigraphs.

Some morals

- The genius, the success, and the limitation of process calculi is their *linguistic character*. This provides an ingenious way of studying processes, information flow, etc. without quite knowing, independently of the particular linguistic setting, what any of these notions are. One could try to say that they are implicitly defined by the calculus. But then the fact that there are so many calculi, potential and actual, does not leave us on very firm ground.

  We lack syntax-independent, *intrinsic* definitions of the fundamental notions of concurrency theory. Net theory and some related approaches (e.g. event structures) still offer the best extant accounts of these issues. But we are still far from home.

  Thus for example consider the issue of *expressiveness*. There are some fragmentary results, but there is no single compelling notion of 'expressive completeness' for a process calculus, or of a 'Church's thesis for concurrency'.

- We must now also acknowledge that we do not have sole ownership of the notions of information, process, etc. Physics and biology are also interested — and they are at our gates! This presents us with a challenge, and perhaps also an opportunity for some new thinking on these issues.

3 New directions: biology, physics or geometry?

A major recent development in process calculi has been their application to biological modelling. This represents perhaps the first substantial example of a trend which, in my view, will form a major part of the future development of our subject: the spreading outwards of ideas developed in Computer Science, of the tool-kits we discussed in Section 1, to other scientific disciplines. Provided there is a real engagement between the CS bio-concurrency community and the biologists, this development has great promise.

However, while biological modelling will surely make new demands on process calculi, and hence lead to new developments (the next 700 biological process calculi?), I don't believe it is likely to lead to foundational advances for the issues we are discussing. Biology's foundational and conceptual structures are, if anything, much more plastic than those of Computer Science — for which, of course, it compensates by the exuberant richness and the sheer concrete reality of the existence proofs which it studies.

There is, perhaps, more prospect for guidance in finding fundamental notions of process, information flow, etc. from the rapidly developing interface between Computer Science and Physics, which has grown up around quantum informatics. We have already discussed how Petri's development of Net theory was influenced by ideas from physics, and indeed provides some of the ingredients of a discrete physics. (One feature conspicuously *lacking* there is an account of the non-local information flows arising from entangled states, which play a key role in quantum informatics. Locality is so plausible to us — and yet, at a fundamental physical level, apparently so wrong!). Meanwhile, there are now some matching developments on the physics side, and a greatly increased interest in discrete models. As one example, the causal sets approach to discrete spacetime of Sorkin et al. [7] is very close in spirit to event structures.

My own recent work with Bob Coecke on a categorical axiomatics for Quantum Mechanics [2,3], adequate for modelling and reasoning about quantum information and computation, is strikingly close in the formal structures used to my earlier work on Interaction Categories [4] — which represented an attempt to find a more intrinsic, syntax-free formulation of concurrency theory; and on Geometry of Interaction [1], which can be seen as capturing a notion of interactive behaviour, in a mathematically rather robust form, which can be used to model the dynamics of logical proof theory and functional computation.

This work admits a striking (and very useful) diagrammatic presentation,

which suggests a link to geometry — and indeed there are solid connections with some of the central ideas relating geometry and physics which have been so prominent in the mathematics of the past 20 years.[3] We note also that, in a rather different style, the geometry of concurrency has been developed by Eric Goubault [5] and others. So, geometry may yet have an important role to play in concurrency theory.

Whither process calculus?

If anything like these speculations comes to pass, I think process calculus will be raised to a new level. It will, perhaps, become truly *the* calculus of a fundamental science of information dynamics.

References

[1] Abramsky, S., Retracing some paths in process algebra. In "Proceedings of the Seventh International Conference in Concurrency Theory" (CONCUR 96), LNCS **1119** (1996), 1–17.

[2] Abramsky, S. and B. Coecke. A categorical semantics of quantum protocols. In "Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science" (2004), 415–425. `arXiv:quant-ph/0402130`

[3] Abramsky, S. and B. Coecke, Abstract physical traces. Theory and Applications of Categories **14** (2005), 111–124.

[4] Abramsky, S., and S. J. Gay and R. Nagarajan. Interaction categories and foundations of typed concurrent programming. "Deductive Program Design: Proceedings of the 1994 Marktoberdorf International Summer School", pp. 35–113. NATO ASI Series F, Springer-Verlag, 1995.

[5] Goubault, E. Some Geometric Perspectives in Concurrency Theory. To appear in Homology, Homotopy and Applications. Available at `http://www.di.ens.fr/Egoubault/GOUBAULTpapers.html` , 2001.

[6] Petri, C.-A., State-Transition Structures in Physics and in Computation. International Journal of Theoretical Physics, **21(12)**, (1982) 979–993.

[7] Sorkin, R., First Steps in Causal Sets. Available at `http://physics.syr.edu/∼sorkin/some.papers/`.

[3] For the afficionado: the diagrammatics of our categories connect with categorical approaches to the Jones polynomial and other topological invariants, which in turn are strongly connected to quantum groups and topological quantum field theories.